University of Michigan Retirement Research Center

Working Paper
WP 2005-103

Estimating Life Cycle Effects of Subjective Survival Probabilities in the Health and Retirement Study

Michael Perry

MRRC Project #: UM05-S1

Estimating Life Cycle Effects of Survival Probabilities in the Health and Retirement Study

Michael Perry
University of Michigan

September 2005

Michigan Retirement Research Center
University of Michigan
P.O. Box 1248
Ann Arbor, MI 48104
http://www.mrrc.isr.umich.edu/
(734) 615-0422

Acknowledgements

This work was supported by a grant from the Social Security Administration through the Michigan Retirement Research Center (Grant # 10-P-98358-5). The findings and conclusions expressed are solely those of the author and do not represent the views of the Social Security Administration, any agency of the Federal government, or the Michigan Retirement Research Center.

Estimating the Life Cycle Effects of Subjective Survival Probabilities in the Health and Retirement Study

Michael Perry

Abstract

This paper attempts to confirm the life-cycle relationship that lower subjective survival probabilities should lead to less positively sloped consumption trajectories. I use the results of six waves of subjective survival probability questions in the HRS to construct an index of survival belief that exploits the panel nature of the data by summarizing all of a respondent's answers to such questions. In conjunction with constructed consumption values from the financial section of the HRS, I test the life-cycle relationship using OLS and Least-Absolute Deviation regression. I find weak evidence that the life-cycle effect of subjective survival probability is significant in a high-cognitive-ability sub-sample of the HRS. Measurement error in the constructed consumption data is problematic.

I. Introduction

People's beliefs about their own life-expectancy have not been extensively studied—mainly due to lack of data. It is not clear that people actually have consistent beliefs about their future chances of survival at any time. Even if they do, measuring them in a meaningful and convincing way is difficult. The Health and Retirement Study has been attempting to obtain such measurements since 1992. In each wave of the survey, respondents are asked what they think the probability is that they will live to be a specific age. Since these questions began being asked in 1992 several studies have examined the question of to what degree respondent's stated probabilities of survival relate to actual survival—this is a question of how informed respondents are about themselves and about human life expectancy. Few studies exist which examine to what extent these subjective survival probabilities affect respondents' decision-making—specifically their financial decision-making—as an economist might predict they would. That is the question under examination.

The life-cycle hypothesis makes a simple prediction about the relationship between a person's perceived survival probability and their consumption: those who think they are more likely to survive will have less consumption growth over time. Simply put, if you expect to live a long time, you will conserve your resources early in life in order to have enough later—this means earlier consumption will be lower than it would have been if you had thought your chances of survival were worse, ceteris paribus. In this way, a higher expected chance of survival should have the same effect as a higher interest rate or a lower degree of impatience.

Most attempts to confirm this relationship in the past have been confounded by lack of data on people's perceived risk of death. Typically, these tests have been based on proxies of life-expectation such as parents' ages of death or life-tables. Testing this implication using life-table mortality rates, Kuehlwein (1993) finds only mixed support for this relationship. Hamermesh (1984) finds some evidence that those who should expect to live longer retire later and consume less. He bases his inference about individuals' expected risk of death on both life-tables and the longevity of individuals' parents. Both of these studies are vulnerable to the criticism that their proxies for individual expectations may not reflect actual individual expectations. This study's main contribution is to examine a similar question using data that may better represent people's actual survival beliefs because it is obtained directly from them.

The Health and Retirement Study (HRS) has elicited subjective life-expectation data from its respondents since the study's inception in 1992 (12 waves of the HRS have been completed—1992-2002, every two years). The questions are of the form "What is the percent chance that you will live to be 75 or more?" (the target age—75 in this case—can vary). Several papers have examined the responses to these questions, and a few broad facts have been established about them: their mean values over respondents are generally close to life-table data; they contain information about respondents' likelihood of dying that the extensive health survey within the study does not capture; respondents' answers to these questions co-vary reasonably with lifestyle indicators of mortality, such as smoking; and these answers appear to be updated as new health information arises between waves of the HRS. However, the responses vary more than life-table data and a large proportion of respondents answer zero, 50 or 100 percent (Hurd, McGarry 1995, 2002).

That these responses broadly agree with life-table statistics, show evidence of updating, and reflect lifestyle choice suggests that respondents answer the questions honestly and take them seriously. That the variance is larger than it would be if respondents took their response from a life-table suggests error in measurement of expectations or that some people are optimistic or pessimistic about their expected longevity as compared to a life-table. The existence of so many answers at zero, 50 and 100 also suggests measurement error or significant rounding, or that respondents have a high degree of uncertainty about the risk of death. Nevertheless, many of the responses seem quite rational and therefore it makes sense to use them, along with other HRS data to test whether respondents' answers affect their financial decision-making. In particular, because some of the subjective survival responses seem sensible and others not, and because the HRS contains other data that indicate mental, it makes sense to test whether there is a subset of respondents for whom the life-expectation data is sensible and who use those expectations to inform their financial decision-making as the life-cycle hypothesis predicts. This follows Hamermesh (1984) again, as he used data from the Terman Study of Gifted Individuals in the hope that if anyone would have the capacity to make decisions using a life-cycle framework, it would be those people who have larger mental capacity.

This study will be a joint examination of three separate questions: how well the life-cycle hypothesis predicts consumption behavior; to what degree HRS respondents (or people in general) take seriously the question of predicting their own survival; and to what degree people's actual beliefs about their mortality risk (whether captured by the survey or not) influence their financial decisions. My main emphasis is on the latter two questions; my life-cycle model will be relatively basic. This point is relevant because the biggest shortcoming of this study will be the lack of direct consumption data from the HRS respondents. If the main goal of this paper were to test the life-cycle implication independent of using the HRS data, then a better research design would probably be to use a survey with a consistent measure of consumption, and then to use a reasonable proxy for survival belief—as in the Kuehlwein and Hamermesh papers mentioned above. I tolerate the lack of consumption data and the problems it introduces because the main goal here is to see whether there is a substantial effect on consumption from people's actual beliefs about their survival probability—not a presumed belief

In order to test the implication that consumption profiles have a more positive slope for respondents who expect to live longer I use the data in the HRS to construct two sets of variables for each respondent. First, I use data on assets, income and capital gains to deduce a value of consumption for each respondent for each time-period between survey interviews. This is problematic, as it introduces substantial measurement error. Second, I use all responses that a respondent has given to any of the subjective survival probability questions over the 12 waves of the study to calculate a linear survival probability profile for each respondent who has answered at least three such questions during the study. This profile is meant to represent the set of survival expectations that both reflects the responses given and is smooth and linear. I then test the hypothesis using ordinary least-squares regression and least-absolute deviation regression.

I find only weak evidence that the relationship exists in the data. This may be due to the measurement error introduced in the process of calculating consumption, or it may be that the relationship actually does not describe respondents' behavior.

In Section One I describe the theory that underlies the life-cycle prediction I test. Section Two contains a description of the data and a description of the calculations used to produce values of consumption and of expected survival probabilities. In section three I test the life-cycle implication that consumption profiles will be more positively sloped for those who have greater subjective survival probabilities.

II. Life-cycle Theory

I use a simplified life-cycle model to produce the implication that I test. The implication that consumption growth should rise with a rise in a person's mortality risk comes directly from the Euler equation of an agent maximizing the sum of additively separable utility over his lifetime (my formulation is borrowed from Kuehlwein):

$$E\left(\frac{U'(C_{t+1})}{U'(C_t)}\cdot p\cdot\frac{1+r}{1+\delta}\right)=1. \qquad (1)$$

Here, p is probability of surviving to the next period, r is the interest rate and δ is the rate of time-preference. Lowering p has the same effect as raising δ or lowering r—it privileges current consumption over future consumption.

This formulation elides the issue of the utility value of a bequest upon dying. In this study I ignore the possible effects of such a value. To what degree people actually behave based on a desire to leave a bequest is an open question that will not be addressed here.

I assume a felicity function $U(C) = \frac{C^{\gamma}}{\gamma}$ (constant relative risk aversion with relative risk aversion parameter 1-γ), and take the logarithm of each side of (1). Also, because my concern is with life-span uncertainty, I assume that the income stream is known. This means that there is no uncertainty about realized consumption in period *t+1*, given that the respondent survives to that period, so I dispense with the expectation operator:

$$\log\left(\frac{C_{it+1}}{C_{it}}\right) = \frac{1}{1-\gamma_i}(\log p_{it} + \log(1+r_{it}) - \log(1+\delta_i)) . \tag{2}$$

The subscript *i* has been added to indicate variables that vary across respondents. Abbreviate $\log\left(\frac{C_{it+1}}{C_{it}}\right)$ by ΔC_{it}. Adding a term, *u*, to account for measurement error in the change in log-consumption gives:

$$\Delta C_{it} = \frac{1}{1-\gamma_i}(\log p_{it} + \log(1+r_{it}) - \log(1+\delta_i)) + u_{it} . \tag{3}$$

As written, this equation should apply only to a single agent making decisions for himself. Analyzing a similar case for a multi-person household in which agents care for each other's well-being requires further assumptions about how those agents interact and make decisions together. For the sake of this analysis I need every data point I can get, so I show results both for singles and for all households. When I analyze a household, I will assume that the household bases its decisions on the well-being of the member who has the highest next-period subjective survival probability. To do this I will use the maximum value of *p* among the people in the household. This assumed decision structure could certainly be replaced with a different one, but this is relatively straightforward and gives the analysis substantially more power as it increases the amount of available data by a factor of four.

The only variables from (3) that I have measured variation in are consumption and subjective survival expectation. One possibility to get variation in the interest rate is to segregate respondents by income-tax bracket thereby dividing people into groups based on after-tax interest rate. I do not do this because it seems unlikely that most respondents are actually that sensitive to what their marginal tax rate is or to what the boundaries between tax brackets are. Therefore, I make the possibly unfounded assumptions that the difference between $\log r_{it}$ and $\log \delta_i$ is distributed randomly in the population given $\log(p_{it})$, and that γ_i is constant (or distributed randomly) throughout the population. This leaves the relationship that I examine:

$$\Delta C_{it} = K + \beta \log p_{it} + u_{it}, \tag{4}$$

substituting $\beta = \frac{1}{1-\gamma}$ and $K = \frac{1}{1-\gamma}\log\left(\frac{1+r}{1+\delta}\right)$.

III. Data

The HRS is a nationally representative panel study of persons over 50 in the United States. Beginning in 1992, respondents were interviewed every two years, covering health, finances, physical and mental capabilities, family structure and relationships and job history. A study called AHEAD (Assets and Health Dynamics of the Oldest Old) began in 1993 and focused on older respondents. In 1996, the AHEAD study merged with the HRS. New cohorts were added to the HRS in 1998 so that the survey would remain representative of those over 50. The last wave of data available for this analysis comes from interviews done in 2002. I employ all the HRS waves, but I do not use AHEAD data that were taken prior to the merger with the HRS.

Table 1 shows descriptive statistics of the population that I will use for this analysis. For each survey wave this population consists of all respondents who answered at least one subjective survival question in that wave. Panel A shows statistics for all those respondents, while Panel B shows the same statistics for those respondents who were single during that time period. P(75) and P(85) refer to the mean values of the probability responses to the subjective survival questions that ask about target ages of 75 and 85 respectively[1]. These statistics are meant simply to make it clear what the population of respondents is like in any particular year. They cannot be used to make accurate inferences about the evolution of households or singles in the HRS population over time because those respondents who have a valid answer to at least one subjective survival question are a highly non-random group. This is due to both self-selection (it takes a certain mental capacity to give a sensible answer to a probability question) and due to survey variation (exactly which sets of respondents have been asked which questions has varied over time in the HRS).

Predictably, the single population has significantly lower assets and income than the overall population. Also, the single population has a higher proportion of females due to women having longer life-spans than men. Finally, the single population generally has a slightly lower subjective survival probability—which could simply reflect the strong relationship between subjective survival probability and wealth that has been noted elsewhere (Hurd, McGarry 1995).

[1] P(85) is a misnomer for the 2000 and 2002 waves because in those waves the target age of the probability question varied based on respondent age. This is described in detail later.

Measuring Consumption

Each wave of the HRS contains detailed questions on household assets (both real and financial), household income (separate from capital gains), and capital gains. The survey does not contain any consistent measure of household consumption. In order to test the implication of survival expectations on consumption profiles, I use the HRS data on assets, income and capital gains to infer a measure of consumption for each respondent for each period between survey interviews.

The basis of the calculation is the relationship:

$$C = I + CG - \Delta A. \tag{5}$$

That is, consumption between two measured points in time equals whatever the household took in—in earned income and capital gains—minus the amount that their asset level grew during that period. It is ambiguous in the HRS whether respondents give pre-tax or post-tax income levels and so there is no way to account for income tax. I do, however subtract property taxes from inferred consumption.

First, I use the HRS income data to estimate household income over the period between survey interviews. I divide the study's constructed household income variable—which estimates total household income in the one-year period prior to the interview—by twelve to get an estimated monthly income and then multiply by the number of months between interviews. This procedure will add measurement error to the extent that actual household income during the period between interviews differs from income during the period just prior to the interview. Additionally, all financial variables in the HRS include imputed values which increase the level of measurement error, but also substantially increase the number of data points available. To exclude the imputed values from this analysis would entail dropping a majority of the available data since almost all respondents require imputation on at least one financial variable.

Second, I use the capital gains section of the survey to estimate capital gains between survey interviews. Respondents are asked whether they have put money in to or taken money out of their various assets. This information, combined with the asset values reported in the earlier and later waves, allows for inference of the respondent's capital gains over the period. This is straightforward except that housing capital gains are not well-measured for respondents who buy or sell a house during the period, so those respondents are dropped.

Finally, I calculate respondents' change in assets between the survey interviews by subtracting the later survey-interview household assets variable from the earlier survey-interview household assets variable. I do not include housing assets on the assumption that people—particularly retired people—do not generally monetize housing assets for the sake of consumption. This assumption is probably alright for the period 1992-2002, but may be less true now as mortgage refinancing for consumption seems to have become much more common.

Adding income and capital gains and subtracting asset growth and property taxes, and then deflating by the CPI-U yields the measure of consumption in 2002 dollars that is used to test the life-cycle prediction.

Using this strategy I have measures of consumption for the periods between the survey waves 1992 and 1994, 1994 and 1996, 1996 and 1998, 1998 and 2000, and 2000 and 2002. These five sets of consumption data can be used to calculate four cross-sections of log-consumption growth, the statistic of interest in the Euler equation. Table 2 shows summary statistics for consumption (in 2002 dollars) and log-consumption growth. Panel A shows consumption measured for all households and consumption measured for households composed of singles. Consumption over the measured periods (approximately two years) rises from $81K in 1992-1994 to $105K in 2000-2002, but with a substantial drop in the prior period to $69K (more on that below). These means do not seem unreasonable for households whose mean yearly income is fairly stable at approximately $60K. Measured consumption for singles is substantially less stable—rising and falling substantially between each period. Panel B shows log-consumption growth measured for all households and for singles. For all households and for singles mean growth is negative in each of the first three periods and is positive in the last period.

Two elements of this table suggest large measurement error. First, in each year a large proportion of households have negative values for this measure of consumption—typically 11-15%. Because actual consumption cannot be negative, these cases are necessarily mis-measured. The proportion of negative cases is a lower bound on the proportion of mis-measured cases in each year. The large number of negative values also explains why the number of cases is significantly lower in Panel B than Panel A—if there is a negative value in either time t or time t+1, then log-consumption growth cannot be measured.

Second is the fact that, for both the whole population and for singles, measured consumption drops substantially for the period 1998-2000 and then rises substantially for the period 2000-2002. This is very likely due to unreported capital gains appearing in the change in asset level. If a respondent had substantial positive capital gains in 1998-2000 (as many did), then did not report them as capital gains and did correctly report their total assets, this would result in measured consumption being biased downwards. The reverse is likely for the period 2000-2002. Poorly measured capital gains are probably not restricted to these time periods—they just show up strongly in these periods because asset values fluctuated substantially.

One final reason to believe that there is a large amount of measurement error in the measure of consumption is the result of a separate study (Perry, 2005, available by request) in which I fit the covariance structure of log-consumption growth to three possible models of the consumption time-series using the generalized method of moments. The three models are: a pure measurement error model in which an individual's log-consumption in each period is an individual specific constant plus a random shock, uncorrelated with any other variable; a model in which log-consumption follows a random walk; and the life-cycle model from equation (4). The equations governing the first two models are

$$\ln c_t = c_i + u_t \tag{6}$$

$$\ln c_{t+1} = \ln c_t + u_t \tag{7},$$

respectively. For both it is assumed that

$$U \equiv \begin{bmatrix} u_1 \\ \vdots \\ u_T \end{bmatrix} \sim N\left(\begin{bmatrix} 0 \\ \vdots \\ 0 \end{bmatrix}, \Sigma \right), \Sigma = \begin{bmatrix} \sigma^2 & 0 & \cdots & 0 \\ 0 & \sigma^2 & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & 0 & \cdots & \sigma^2 \end{bmatrix}.$$

Using a χ^2 goodness-of-fit test, I am able to reject the random-walk and life-cycle models and unable to reject the measurement error model. While it is possible that the measurement error model is describing an actual consumption process, rather than measurement error, this seems unlikely. While this result does indicate that the overall model fits the data poorly, I persist in order to see whether there is a positive relationship between log-consumption growth and *p*.

Change in log-consumption is the dependent variable in my analysis. First-differencing of consumption will exacerbate measurement error in levels of consumption. Regression coefficient values should not be biased due to random measurement error in the dependent variable, but they will be measured less-precisely. Moreover, I have no evidence that measurement error in this measure of consumption is random—to the extent that it is not, coefficient values could be biased.

The HRS does provide a few variables that can be used to corroborate my deduced consumption values. In 1996 and 1998 the survey asked each household what their total spending—including all debt payments, utility bills, rent, transportation, entertainment, food, clothes and any other expenses—was in the previous month. Also, in 2002, the survey asked three food consumption questions: how much did the household spend in the past week on all food; how much did it spend having food delivered; and how much did it spend eating out. The left side of table 3 shows mean values for these measures and for my deduced consumption values measured on a monthly basis for the same time period. The levels of total monthly consumption from the HRS survey are substantially lower than my calculated values in both 1996 and 1998. Also included are mean values of inferred consumption with negative values removed—this increases the difference between the HRS measure and my measure. It is questionable how accurate a respondent is likely to be in making a fast estimate of monthly spending, so there is no guarantee that the HRS measure is very good. Indeed, given that the average income levels for this population in 1996 and 1998 are $64K and $73K respectively, the levels of spending implied by the HRS measure seem quite low and would imply a very high average savings rate. One possible, partial explanation for the large difference between the values reported by respondents and the calculated values is that I have not accounted for income tax. If respondents generally report pre-tax income, then my calculation will count their taxes as consumption. It seems likely that few respondents would include income tax in their response to the 1996 and 1998 HRS consumption question.

The right side of table 3 shows correlation coefficients and respective significance levels between the HRS measures and my inferred levels of consumption for the relevant time periods. For both HRS consumption measures, the correlation is substantially higher when the negative cases are removed from the deduced consumption numbers. This is unsurprising as those cases almost certainly represent particularly egregious cases of measurement error. Furthermore, it is encouraging that the inferred consumption shows a substantial correlation with the HRS measures of consumption, given the disparate measurement techniques and the likely presence of substantial measurement error in both. Interestingly, of the food consumption measures in 2002, only the measurement of what a family spends eating at restaurants is significantly correlated with my inferred consumption measure. Also interesting, though I have not shown it, is that these measures of food consumption correlate very little with each other.

Table 3 shows that despite its weaknesses, my inference about household consumption do match up to a substantial degree with the limited information the HRS survey provides about actual household consumption.

Measuring Subjective Survival Expectations

In the Euler equation, *p* represents the agent's subjective assessment of his probability of living to the next period. The HRS provides answers to questions of the form "What is the percent chance that you will live to be 75 or more?" These questions are asked twice in each survey wave with different target ages, although some respondents may only be asked once or not at all. From 1992 to 1998 respondents were asked the questions with 75 as a target age, and then with 85 as a target age. In 2000 and 2002, the first question remains the same and the second question has a target age that varies from 80 to 100 in five year increments depending on the age of the respondent (the target for anyone under 70 was 80, for those 70-74 it was 85 and so on).

Figure 1 shows a histogram of the answers to the "75" question in 2002. Figure 1 is representative of responses in any given year in that a very large proportion of people answer 100 or 50 in addition to substantial fractions at 75, 80, and zero.

Table 4 shows statistics of attributes of people who gave certain answers to P(85) in each year of the survey. Panel A shows the differences between all respondents and those who gave an answer of 50 or more and those who answered less than 50. Those who give lower answers are more likely to be male—a fact that squares well with actual mortality data. Additionally, those who give lower answers have less education, fewer assets and less income than those who gave higher answers.

Panel B shows the same statistics for those who gave answers of exactly zero, 50 and 100. Because such a high proportion of respondents give these answers and because in the case of zero and 100, they are not very sensible answers, it is worth checking whether there is something noticeably different about these respondents. In contrast to the result from panel A, those who answer 100 have less education, fewer assets and less income than those who answer 50. Those who answered zero have less still. Those who answered 50 look essentially the same as the rest of the population. This suggests that answers of zero and 100 may be more a sign of poor understanding of the question than of optimism or pessimism. This conclusion will be used in developing a measure of expected next-period survival probability.

A response to one of these questions does not imply directly any particular value of the respondent's expected chance of living to any particular date other than the target age. In order to use the survey responses to calculate a value of *p* (in the life-cycle model the probability of living to the next period; in this analysis the probability of living through the next period of measured consumption) for each respondent, some assumptions are necessary.

For each respondent, I assume that a response to the question "What is the percent chance that you will live to be 75 or more?" implies a belief over all the conditional probabilities of surviving one year into the future (that is, for example, the probability of surviving to age 63 given that the respondent has survived to age 62) for each year from the respondent's current age up to the age of 75. Assuming that these probabilities exist, the response to the question is just the product of all conditional probabilities from the respondent's age to the target age:

$$R = \prod_{i=A+1}^{i=T} \rho_i \,. \qquad (8)$$

Here, R is the survey response, A is the age of the respondent, T is the target age and ρ_i is the probability of surviving to age i, given that the respondent has survived to age i-1. In order to calculate values of ρ_i, it is necessary to assume something about how respondents' beliefs change over time. I assume that respondents recognize that their conditional survival probabilities fall somewhat each year that they age[2]. Over the relatively short time of a decade, actual life-table survival probabilities decline approximately linearly. For this reason and for simplicity, I assume that ρ_i declines by a constant amount each year. This simplifies the expression in (8) to

$$R = \prod_{i=1}^{i=T-A} (\rho_{A+1} - (i-1)m), \tag{9}$$

where *m* is the amount by which survival probabilities decrease each year and *H=T-A*. Taking logs of both sides gives

$$\ln R = \sum_{i=1}^{i=H} \ln(\rho_{A+1} - (i-1)m). \tag{10}$$

Then, set $\rho_{A+1} = 1 - r$. Actual mortality rates, even for people in their eighties are typically below 0.1—meaning that in actual outcome, survival probabilities are quite close to one for any given year. I assume that respondents' beliefs conform well-enough to actual outcomes that I can use the approximation $\ln(1-x) \cong -x$, for small *x,* in equation (10). This yields

$$\ln R = \sum_{i=1}^{i=H} (-r - (i-1)m) \tag{11}$$

$$\ln R = -rH - m\frac{(H-1)(H)}{2}, \tag{12}$$

[2] This assumption may be reasonable for those respondents who gain no new and significant information about their life-expectancy during the relevant time period. It is almost certainly not reasonable for respondents who do receive such information by, for example, suffering a major health shock such as a stroke.

where equation (12) follows from standard summation results.

This is the equation I use to describe the relationship between a response and a respondent's beliefs in the first wave of the survey (1992). Because other responses occur at different times and for different target ages, there is variation in the values of *H* and therefore in the multipliers of *r* and *m*. For example, if equation (12) represents the relationship between beliefs and response for the question with target age 75, asked in 1992, then the same relationship for target age 80, asked in 2000 of the same respondent looks like:

$$\ln R_{2000,80} = -rH_{2000,80} - m\left(\frac{(H_{2000,80}-1)(H_{2000,80})}{2} + 8\right) \tag{13}$$

where $H_{2000,80} = H - 8 + 5 = H - 3$. The difference in *H* is due to the respondent's age having advanced eight years between surveys and the target age increasing by five years. The addition of 8 to the multiplier on *m* is due to all conditional survival probabilities having declined by *8m* as the respondent aged during the time between 1992 and 2000.

Using these relationships, I have a vector of responses, *R*, and a matrix of multipliers for *r* and *m*, *X,* for each respondent. This allows me to estimate the regression

$$\ln R = X'\begin{pmatrix} r \\ m \end{pmatrix} + \varepsilon \tag{14}$$

separately for each respondent, thereby giving values of *r* and *m* for each respondent who has answered at least three subjective survival questions. In this formulation, *r* is the respondent's perceived risk of death in the first year (the first year is set to 1992 for all respondents since that is the first year of survey data for any respondents), and, again, *m* is the yearly increase in risk of death. Using these numbers, I calculate a respondent's perceived risk of death in year x as $risk_x = r + (x-1992)m$, or equivalently, I calculate their perceived probability of survival during year x as $p_x = 1 - risk_x$.

The above explanation ignores the issue that when a respondent answers zero, it is impossible to take a log and use that response in the calculation. I try two strategies to deal with this issue and test which seems to work better. First, I exclude all responses of zero from the calculations. Above, we established that those who answer 100 seem to be similar to those who answer zero, and additionally, they seem to be the same sort of unlikely answer to a probability question—perhaps due to misunderstanding. For that reason, when I exclude the zeros I also exclude the 100s. For the second strategy, instead of excluding the zeros and 100s, I replace the zeroes with the value 0.00001, which can be logged, and I replace the 100s (really 1s since everything is converted to fractions) with 0.99999.

In each case, I use the values of r and m generated for each respondent to calculate the respondent's perceived probability of survival during any year. These predicted yearly subjective survival values can be multiplied together as in (8) to produce predicted responses to any of the subjective survival questions on the survey. To test my two strategies, I regress the actual responses on the predicted responses. The results are shown in Table 4. Strategy 1 drops responses of zero or 100, strategy 2 replaces them. The first two sets of R^2 values are for regressions over the same responses. The last set is for strategy 2 used to predict for all responses for which it is possible to do so. The number of possible cases using strategy 2 is larger because dropping responses in strategy 1 necessarily means reducing some respondents to below the three-response level necessary for prediction. Strategy 1 produces a better set of predicted responses in all cases. This could be because answers of zero or 100 are more likely to reflect confusion than information about held beliefs.

Perhaps needless to say, I do not hypothesize that any respondent has set beliefs about his or her conditional probability of surviving during any particular year. It would be claiming too much to say that the HRS questions evoke anything more than a general impression of survival probability from most respondents (the exception perhaps being any professional actuaries surveyed). The scheme I propose for integrating all of a respondent's answers is intended to be a fairly straightforward way of approximating what a respondent's well-articulated beliefs might look like if they were forced to develop them in a rigorous way and if they had some consistency over time. Therefore, the charge could be leveled that I have invented an index with a dubious epistemic nature. My only response is that I see no other simple strategy for incorporating all of a respondent's answers that is not at least as questionable. It may well be that questions like those on the HRS are simply not sophisticated enough to use in testing life-cycle models.

IV. Results

To test whether the life-cycle prediction holds I first check to see whether it holds in a coarse or general way before I attempt to estimate a precise effect. As noted above, six waves of HRS data yield five periods of consumption data, which in turn can be used to produce four data points per respondent of log-consumption growth. I calculate predicted subjective survival beliefs for each respondent corresponding to their belief that they will survive from the time of the HRS interview all the way through the next period of calculated consumption. For example, I have consumption measured for the periods 1992-1994 and 1994-1996. This allows me to calculate log-consumption growth for the period 1992/1994-1994/1996. Then, the relevant survival belief to juxtapose with log-consumption growth from 1992/1994-1994/1996 is the respondent's belief that they will live from the survey interview date in 1994 all the way *through* their next period of measured consumption—which I assume to be the end of 1996. Therefore I calculate the respondent's subjective belief that he or she will survive through the years 1995 and 1996. I do this calculation for each respondent, for each measurement of log-consumption growth. I then aggregate the four cross-sections of log-consumption growth data and subjective survival probability data into one dataset.

Figures 2 and 3 shows a scatter plot of the basic data: log-consumption growth vs. log-subjective survival belief for households and for singles. There is no clear relationship, but the existence of some very low subjective survival beliefs distorts the abscissa. Figures 4 and 5 fix this by showing the same data, with the *x*-axis truncated below -0.3. Again, there is no clear relationship between the sets of points.

Figures 6 through 9 clarify the relationship somewhat by plotting the mean log-consumption growth for each decile of log-survival belief (the higher deciles indicate a higher subjective survival probability). In figures 6 and 7 there does appear to be a noisy, but positive relationship between the variables as the Euler equation predicts. However, figures 8 and 9—in which the same data is graphed, but including bars showing one standard deviation on each side of the mean—show that the log-consumption growth data has so much random variation in it that we cannot have any confidence in the relationship that figures 6 and 7 show.

Regressions of log-consumption growth on *log p* show similar results: the sign of the relationship is correct, but the significance is too low to have any confidence in the result. Table 6 shows the results of OLS regressions of log-consumption growth on *log p* for the entire sample. The first line of table 6 shows the results when the only independent variable is *log p*. The second line includes *log p*, years of education, an indicator for white race, and the decile of asset level of the household as regressors (using asset deciles instead of asset levels gives very similar results and gives the coefficient much more convenient values for reporting). In both of the first two lines, the sign on *log p*'s coefficient is positive, but not significant at standard significance levels.

The next two lines show results of the same regressions for a sub-sample selected based on cognitive ability. The HRS includes a sequence in which respondents are read a list of 20 words and asked to recall as many as possible. Then, a few minutes later, they are again asked to recall as many as possible. The HRS also includes a set of questions in which the respondent is asked "what is 100 minus seven?", "and seven from that?", etc. With the question repeated three more times. I have taken the total number of words a respondent remembered after being asked each time and the total number of correct answers they gave to the "serial 7s" questions to produce a word recall score and a serial 7s score. In addition, we know the level of education for the respondents. I select a sub-sample of respondents who are above the median on all three cognitive ability measures.

The reasoning behind this is that the life-cycle hypothesis applies to rational agents maximizing lifetime utility. Presumably, it takes a good deal of cognitive sophistication to do this; so perhaps the life-cycle result that we are testing holds for those who are more capable mentally. The results of the OLS regressions do not bear out this hypothesis. Again, the coefficients on *log p* are positive, but not significant.

In the second half of the table I perform the same regressions using least-absolute-deviation (or median) regression instead of least-squares. I do this to minimize the impact of outliers—specifically outliers that may be the result of measurement error due to the calculations I used to produce values for consumption and for *log p*. The results for the entire sample are in the first two lines of the second half of table 6. Here, again, I use *log p* alone as a regressor, and then log p along with education, race and asset level (in deciles). The results for *log p* are similar to the OLS regressions: the coefficient on *log p* is positive but insignificant.

The results for median-regression restricted to the high-cognitive-ability sample—the last two lines of table 6—have positive and significant coefficients on *log p* with t-statistics of 2.06 and 2.31 for the unconditional and conditional cases, respectively. I interpret this as weak evidence for the life-cycle hypothesis—weak because it only shows up in these rather esoteric specifications.

One other aspect of the results to note is that the coefficients on *log p* are considerably larger for the high-cognitive ability sub-sample than for the whole sample in all specifications. I offer two possible interpretations of this. If the results actually indicate life-cycle behavior in the whole population, then the coefficient on *log p* can be interpreted as $\frac{1}{1-\gamma}$ where $1-\gamma$ is the relative risk-aversion parameter. Therefore a higher coefficient on *log p* for a population could be interpreted as a lower value of $1-\gamma$ for that population—i.e. that population is less risk-averse at the same level of consumption. For the total population, an approximate value for $\frac{1}{1-\gamma}$ is 0.15 and for the high-cognitive ability sub-sample is 0.75. These yield relative risk-aversion parameters of approximately 6.7 and 1.3, respectively. Hall (1988) reports results corresponding to values for $\frac{1}{1-\gamma}$ ranging from -0.4 to 0.98. The values from this analysis do not seem far out of line with those results. Also, Hurd (1989), in a study on the effects of mortality risk on consumption and wealth in the Longitudinal Retirement History Survey, estimates a relative risk aversion parameter of 1.12—not out of line with the value obtained here for the high-cognitive ability sub-sample.

The second possibility (which I consider at least as likely), also conditional on some life-cycle behavior actually existing in the population, is that with random measurement error in *log p*, the coefficient on *log p* will be biased towards zero in a standard regression. It could very well be that the cognitively-able portion of the population is better able to produce a survey response that reflects their actual subjective beliefs. This would mean less measurement error in the calculated values of *p,* and therefore coefficients with greater absolute value. If this were the case, then the coefficient values for the high-cognitive ability group might better reflect the true parameter values in the population.

Table 7 repeats the same regressions for single population. These regressions have coefficient values similar to table 6. The major difference is that the sample size is much smaller and hence, standard errors are much larger. Indeed, the t-statistics for *log p* are never such that we consider the relationship significant for singles. The feature that the coefficients on *log p* are higher for the high-cognitive ability group remains.

As noted in footnote 2, my predicted values for subjective survival beliefs should really only reflect the beliefs of those respondents who do not significantly revise their survival beliefs during the decade over which these measurements are taken. An indicator in the HRS data for such a revision would be a major health event such as the diagnosis of cancer. I have not dropped those respondents who suffer a major health event because, as is apparent in the regression results for singles, I need a very large number of data points to get any significant results. A profitable revision to this work would be to develop a more sophisticated method of incorporating all the information in HRS survey on respondents' subjective survival beliefs—a method that could adequately handle major revisions to beliefs. This task should become easier as more waves of survey data become available. For the time being, I can only hope that the effects of revisions do not distort my measurements of survival beliefs too much.

One final major point about the regression results is that in no case is the R^2 value above 1%. If I have shown weak evidence for life-cycle behavior based on subjective survival beliefs, then that relationship explains almost none of the variation in the data. One explanation of that may be measurement error in consumption values. A major shortcoming in the study of subjective survival beliefs on financial decision-making is the lack of joint data on survival beliefs and consumption. Each new wave of the HRS will add significant power to this study, though. As this is being written, the 2004 wave of HRS data is forthcoming in the next six months, which should allow for a better understanding of whether a real relationship exists between subjective survival beliefs and consumption trajectories.

V. Conclusions

In this paper, I have examined constructed consumption and subjective survival probability data to test whether those respondents who believe they are more likely to survive have higher growth in consumption. My constructed subjective survival probability data uses a novel process to integrate all subjective survival responses for any given respondent into one survival curve. These curves alone can explain a large amount of the variation in subjective survival response and they allow for use in predicting subjective survival probability over short periods—which is necessary for estimating life-cycle effects.

My constructed consumption data contains a large amount of measurement error. This makes finding any life-cycle effects difficult. I find weak evidence that such a relationship exists in the data for those respondents who are of higher cognitive ability. I also find some evidence that high-cognitive-ability respondents either have a different group mean of relative risk aversion than the general HRS population or that they are better at articulating their subjective survival beliefs. The value of this study should increase significantly as HRS 2004 data becomes available in the next six months.

References

Hall, R. (1988): "Intertemporal Substitution in Consumption," *The Journal of Political Economy*, 96, 339-357.

Hamermesh, D. (1984): "Life-cycle Effects on Consumption and Retirement," *Journal of Labor Economics*," 3, 353-370.

Hurd, M. (1989): "Mortality Risk and Bequests," *Econometrica*, 57, 779-813.

Hurd, M. and McGarry, K. (1995): "Evaluation of the Subjective Probabilities of Survival in the Health and Retirement Study," *The Journal of Human Resources*, 30, S268-S292.

Hurd, M. and McGarry, K. (2002): "The Predictive Validity of Subjective Probabilities of Survival," *The Economic Journal*, 112, 966-985.

Kuehlwein, M. (1993): "Life-Cycle and Altruistic Theories of Saving with Lifetime Uncertainty," *The Review of Economics and Statistics*, 75, 38-47.

Perry, M. (2005): "On the Covariance Structure of Changes in Consumption," unpublished.

Table 1: Descriptive Statistics for HRS Respondents

A: All Respondents	1992	1994	1996	1998	2000	2002
N	11658	10237	9648	9417	15667	14498
mean age	56	57	59	57	68	67
%male	45	43	43	38	42	41
%white	73	75	76	76	78	77
mean p(75)	65	64	65	65	67	66
mean p(85)	44	44	48	46	51	50
mean assets (2002 $K)	174	208	246	255	290	258
mean annual income (2002 $K)	61	63	64	73	58	58

B: Single Respondents	1992	1994	1996	1998	2000	2002
N	2374	2970	3479	2411	2401	2285
mean age	56	58	60	63	64	66
%male	31	35	38	40	41	41
%white	56	59	60	62	64	64
mean p(75)	62	62	62	64	64	66
mean p(85)	43	46	49	48	54	55
mean assets (2002 $K)	68	86	110	118	140	116
mean annual income (2002 $K)	28	27	31	33	32	32

Table 2: Measured Consumption and Log-Consumption Growth

A: Consumption

All Households	N	mean (2002 $K)	stand. dev. (2002 $K)	%negative
1992-1994	6888	81	242	15
1994-1996	6343	82	276	13
1996-1998	6153	91	463	12
1998-2000	12565	69	393	13
2000-2002	11532	105	338	11

Singles	N	mean	stand. dev.	%negative
1992-1994	2040	33	121	19
1994-1996	1952	46	295	15
1996-1998	2038	59	311	12
1998-2000	5698	37	273	14
2000-2002	5506	66	301	13

B: Log-Consumption Growth

All Households	N	mean	stand. dev.
92/94-94/96	4586	-0.004	1.26
94/96-96/98	4430	-0.070	1.23
96/98-98/00	4224	-0.055	1.24
98/00-00/02	8298	0.060	1.25

Singles	N	mean	stand. dev.
92/94-94/96	1188	-0.075	1.29
94/96-96/98	1282	-0.021	1.25
96/98-98/00	1327	-0.045	1.21
98/00-00/02	3560	0.049	1.22

Table 3: Comparison of HRS Consumption Measures with Inferred Consumption

	mean (2002 dollars)
1996 (survey)	1722
1998 (survey)	1959
1996 (inferred)	3682
1998 (inferred)	4398
1996 (inferred, ≥0)	5438
1998 (inferred, ≥0)	6382
All food (2002 weekly)	86
Restaurants	26
Delivered	1

	correlations	
	1996 (inferred)	1996 (inferred, ≥0)
1996 (survey)	0.14	0.30
	1998 (inferred)	1998 (inferred, ≥0)
1998 (survey)	0	0.33
	2002 (inferred)	2002 (inferred, ≥0)
All food	0	0
Restaurants	0.14	0.18
Delivered	0	0

Table 4A: Summary Statistics for Respondents to P(85)

	N	% Male	P(85)	Age	Years Education	Assets (nominal $K)	Income (nominal $K)
all							
1992	11740	45%	44	56	12.1	136	48
1994	9524	44%	44	58	12.4	179	54
1996	8967	43%	48	59	12.5	223	58
1998	8795	38%	46	58	12.8	244	68
2000	15227	42%	51	67	12.5	284	57
2002	14146	41%	50	68	12.6	262	58
mean	11400	42%	47	61	12.5	221	57
live to 85≥50							
1992	5846	40%	71	55	12.4	154	52
1994	5003	40%	69	57	12.6	189	56
1996	5082	40%	71	59	12.7	234	61
1998	4820	33%	70	58	13.0	272	73
2000	9713	40%	70	66	12.8	323	62
2002	8788	39%	70	67	12.8	295	65
mean	6542	39%	70	60	12.7	245	62
live to 85<50							
1992	5894	49%	16	56	11.8	119	45
1994	4567	47%	18	58	12.2	169	51
1996	3792	48%	17	59	12.3	208	54
1998	3975	44%	18	58	12.6	209	62
2000	5514	46%	16	68	12.0	214	47
2002	5358	46%	16	70	12.2	209	48
mean	4850	47%	17	62	12.2	188	51

Table 4B: Summary Statistics for Respondents to P(85)

	N	% Male	Age	Years Education	Assets (nominal $K)	Income (nominal $K)
answered 0						
1992	2165	50%	57	10.7	89	36
1994	1035	41%	59	11.3	136	42
1996	977	46%	60	11.3	129	42
1998	852	42%	58	11.8	118	46
2000	1441	45%	73	11.0	145	33
2002	1430	42%	74	11.5	154	33
mean	1317	44%	64	11.3	129	39
answered 100						
1992	1136	42%	56	11.8	121	42
1994	803	42%	58	11.8	148	49
1996	1044	41%	60	11.9	163	49
1998	842	30%	58	12.3	191	59
2000	1633	38%	66	11.9	261	51
2002	1480	36%	67	12.0	240	55
mean	1156	38%	61	12.0	187	51
answered 50						
1992	1879	43%	56	12.2	159	51
1994	2049	41%	57	12.4	186	54
1996	1935	41%	59	12.5	194	55
1998	1934	35%	57	12.9	297	71
2000	3692	43%	67	12.6	290	56
2002	3380	43%	68	12.6	269	58
mean	2478	41%	61	12.5	233	58

Table 4C: Summary Statistics for Respondents to P(85) (singles only)

	N	% Male	P(85)	Age	Years Education	Assets (nominal $K)	Income (nominal $K)
all							
1992	2259	31%	43	56	11.9	54	23
1994	1903	29%	46	58	12.2	83	24
1996	1904	28%	49	59	12.3	110	30
1998	1990	28%	48	59	12.6	114	35
2000	4768	25%	48	70	11.9	152	30
2002	4635	25%	46	71	12.1	138	30
mean	2910	28%	47	62	12.2	108	29
P(85)≥50							
1992	1103	27%	73	56	12.2	61	25
1994	1036	25%	70	58	12.4	86	26
1996	1107	24%	73	60	12.4	121	37
1998	1141	25%	72	59	12.8	117	39
2000	2871	23%	70	69	12.2	168	33
2002	2671	23%	71	69	12.3	132	33
mean	1655	25%	71	62	12.4	114	32
P(85)<50							
1992	1156	35.0%	14	56	11.5	47	21
1994	867	32.3%	17	58	11.9	79	22
1996	797	33.8%	16	59	12	93	27
1998	849	32.9%	16	59	12.3	109	31
2000	1897	26.5%	13	72	11.6	128	25
2002	1964	26.8%	13	73	11.9	130	27
mean	1255	31%	15	63	11.9	98	25

Table 4D: Summary Statistics for Respondents to P(85) (singles only)

	N	% Male	Age	Years Education	Assets (nominal $K)	Income (nominal $K)
answered 0						
1992	496	38.9%	56	10.6	31	16
1994	234	23.9%	58	11.1	48	16
1996	218	33.9%	60	11.2	69	21
1998	207	30.0%	59	11.6	68	21
2000	641	23.4%	76	10.8	88	20
2002	702	22.2%	77	11.4	121	22
mean	416	28.7%	64	11.1	71	20
answered 100						
1992	253	23.3%	56	12.1	25	18
1994	199	24.6%	58	11.6	58	19
1996	276	23.9%	60	11.7	88	23
1998	256	22.3%	59	12.2	58	26
2000	547	23.0%	69	11.4	121	26
2002	524	22.7%	69	11.6	73	28
mean	343	23.3%	62	11.8	71	23
answered 50						
1992	329	31.3%	56	12.0	92	26
1994	413	26.4%	58	12.4	106	31
1996	394	26.6%	60	12.2	87	29
1998	438	26.7%	59	12.8	153	47
2000	1174	22.6%	71	12.1	149	29
2002	1087	25.1%	71	12.1	129	31
mean	639	26.5%	62	12.3	119	32

Table 5: R-squared values for Regressions of Actual on Predicted Responses

		Strategy 1	Strategy 2		Strategy 2 on all eligible cases	
		R2	R2	N	R2	N
1992	P(75)	69	46	6396	49	8310
	P(85)	50	35	6561	51	8382
1994	P(75)	60	39	6390	44	8446
	P(85)	55	38	6828	49	8141
1996	P(75)	51	35	5601	44	8074
	P(85)	61	44	6369	52	7879
1998	P(75)	50	36	6139	41	8320
	P(85)	70	52	6246	56	7400
2000	P(75)	50	35	5761	41	7585
	P(85)	45	37	7314	44	9372
2002	P(75)	44	31	4776	38	6288
	P(85)	54	37	7534	45	9086

Table 6: Regression Results for Log-Consumption Growth on log-*p*

	constant	p	education	white	asset decile	N	R-squared
Coefficient	0.01	0.11				14886	<1%
Std. Error	(0.01)	(0.13)					
Coefficient	-0.04	0.11	0.005	0.04	-0.010	14191	<1%
Std. Error	(0.05)	(0.14)	(0.004)	(0.03)	(0.004)		
	Top Half of Word Recall, Serial 7s and Education						
	constant	p	education	white	asset decile		
Coefficient	0.10	0.74				1065	<1%
Std. Error	(0.05)	(0.56)					
Coefficient	-0.06	0.70	0.001	0.17	-0.0016	1065	<1%
Std. Error	(0.41)	(0.57)	(0.027)	(0.10)	(0.017)		
	Median Regression Log-Consumption Growth on log-p						
	constant	p	education	white	asset decile		
Coefficient	0.02	0.15				14886	<1%
Std. Error	(0.01)	(0.10)					
Coefficient	-0.02	0.19	0.01	0.01	-0.01	14191	<1%
Std. Error	(0.04)	(012)	(0.003)	(0.021)	(0.004)		
	Top Half of Word Recall, Serial 7s and Education						
	constant	p	education	white	asset decile		
Coefficient	0.15	0.78				1065	<1%
Std. Error	(0.03)	(0.38)					
Coefficient	0.00	0.77	0.00	0.16	0.00	1065	<1%
Std. Error	(0.25)	(0.33)	(0.02)	(0.06)	(0.01)		

Table 7: Regression Results for Log-Consumption Growth on log-*p (singles only)*

	constant	p	education	white	asset decile	N	R-squared
Coefficient	0.00	0.15				3909	<1%
Std. Error	(0.03)	(0.21)					
Coefficient	0.01	0.14	0.01	0.01	-0.02	3740	<1%
Std. Error	(0.09)	(0.22)	(0.01)	(0.04)	(0.01)		
	Top Half of Word Recall, Serial 7s and Education						
	constant	p	education	white	asset decile		
Coefficient	0.01	0.52				632	<1%
Std. Error	(0.07)	(0.66)					
Coefficient	-0.30	0.80	0.02	0.24	-0.02	590	<1%
Std. Error	(0.58)	(0.69)	(0.04)	(0.13)	0.02		
	Median Regression Log-Consumption Growth on log-p						
	constant	p	education	white	asset decile		
Coefficient	0.02	0.11				3909	<1%
Std. Error	(0.02)	(0.15)					
Coefficient	0.05	0.20	0.01	-0.01	-0.03	3740	<1%
Std. Error	(0.06)	(0.15)	(0.01)	(0.03)	(0.01)		
	Top Half of Word Recall, Serial 7s and Education						
	constant	p	education	white	asset decile		
Coefficient	0.05	0.69				632	<1%
Std. Error	(0.06)	(0.61)					
Coefficient	-0.34	0.74	0.03	0.13	-0.02	590	<1%
Std. Error	(0.61)	(0.71)	(0.04)	(0.14)	(0.02)		

Figure 1: Histogram of Responses to P(75) in 2002

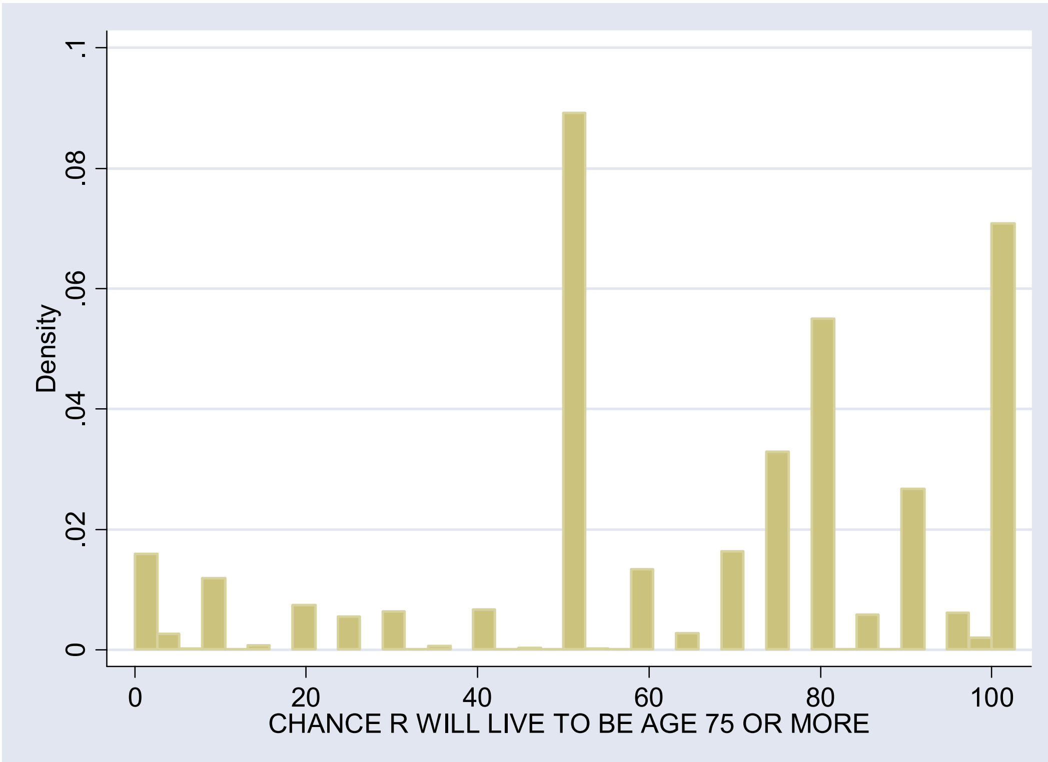

Figure 2: Log-Consumption Growth vs. Log-Survival Probability

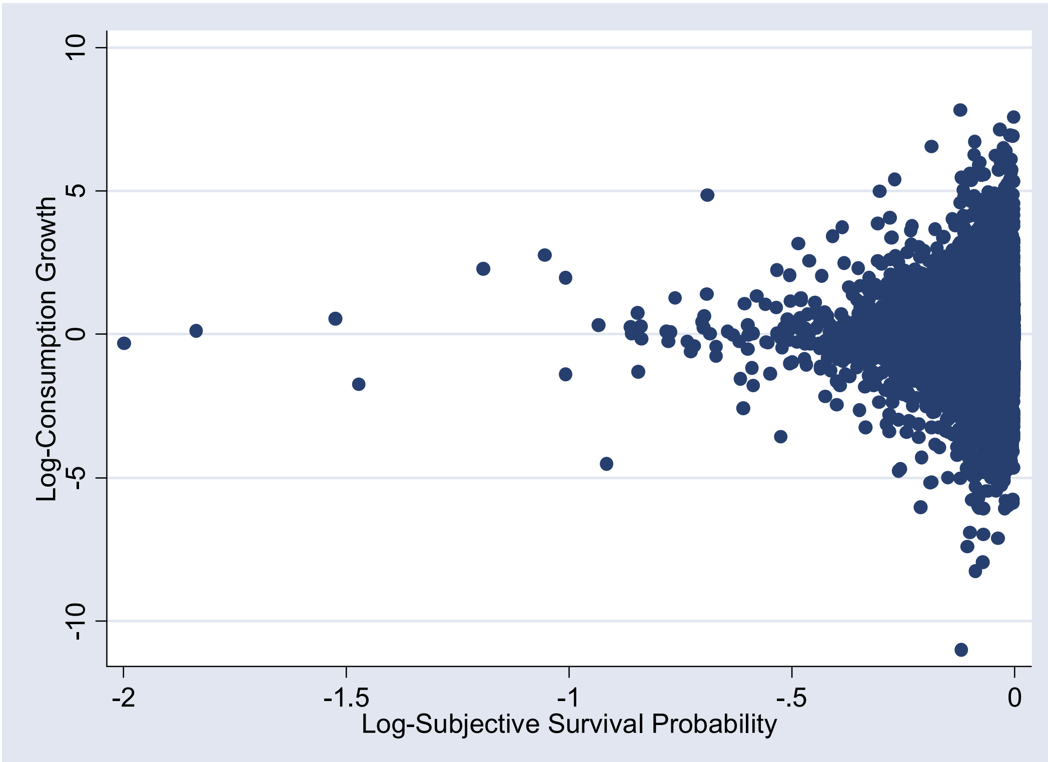

Figure 3: Log-Consumption Growth vs. Log-Survival Probability (Singles Only)

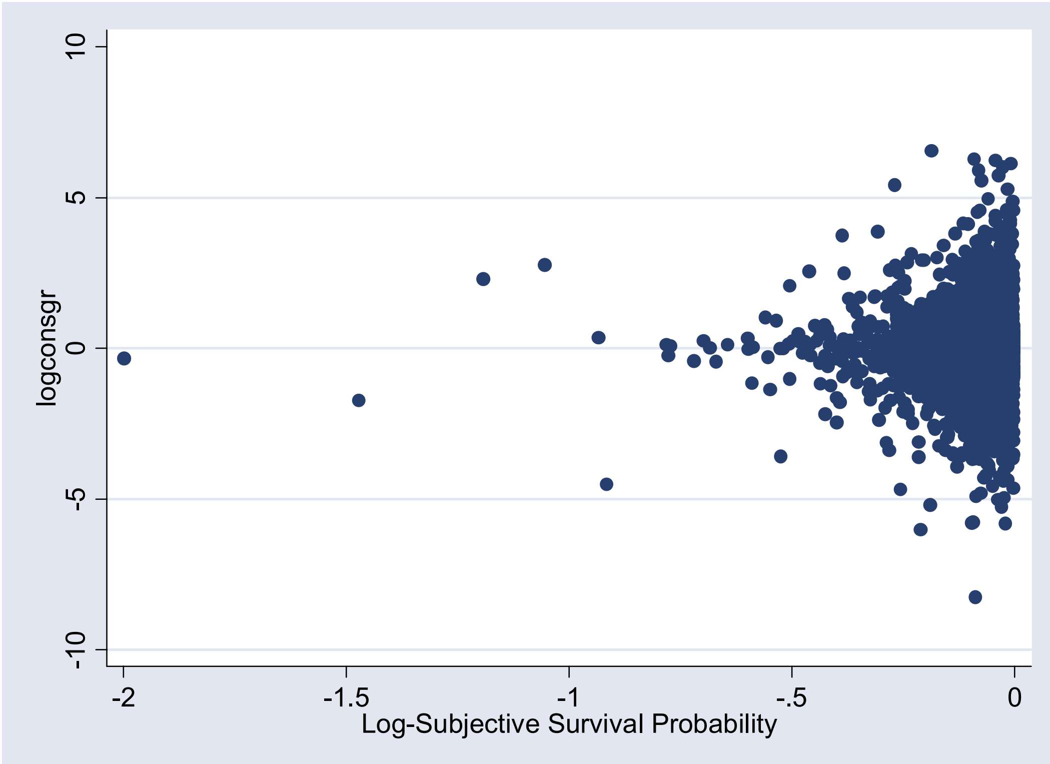

Figure 4: Log-Consumption Growth vs. Log-Survival Probability (Truncated Scale)

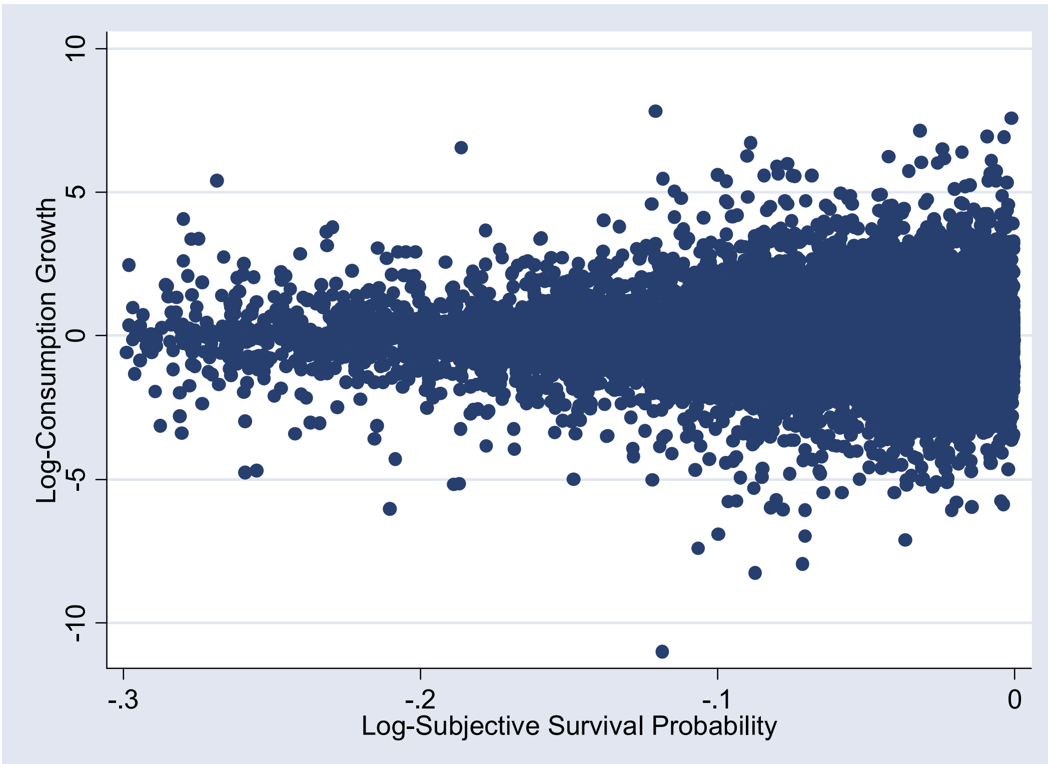

Figure 5: Log-Consumption Growth vs. Log-Survival Probability (Singles Only, Truncated Scale)

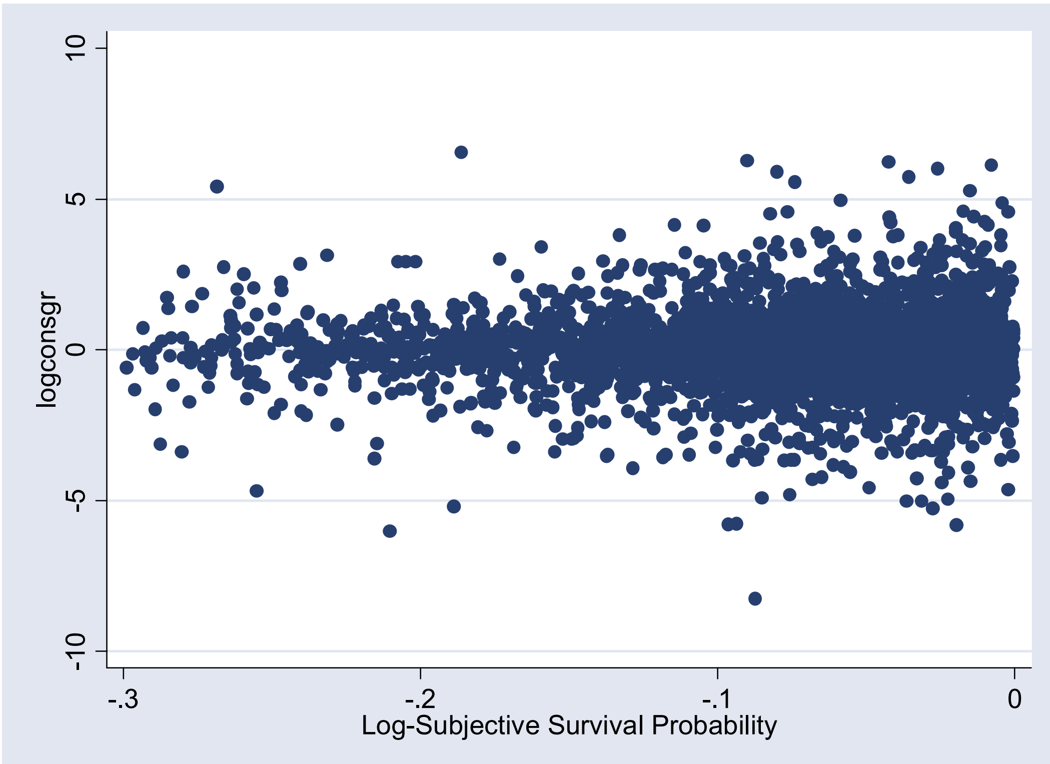

Figure 6: Mean Log-Consumption Growth by Decile of Log-Survival Probability

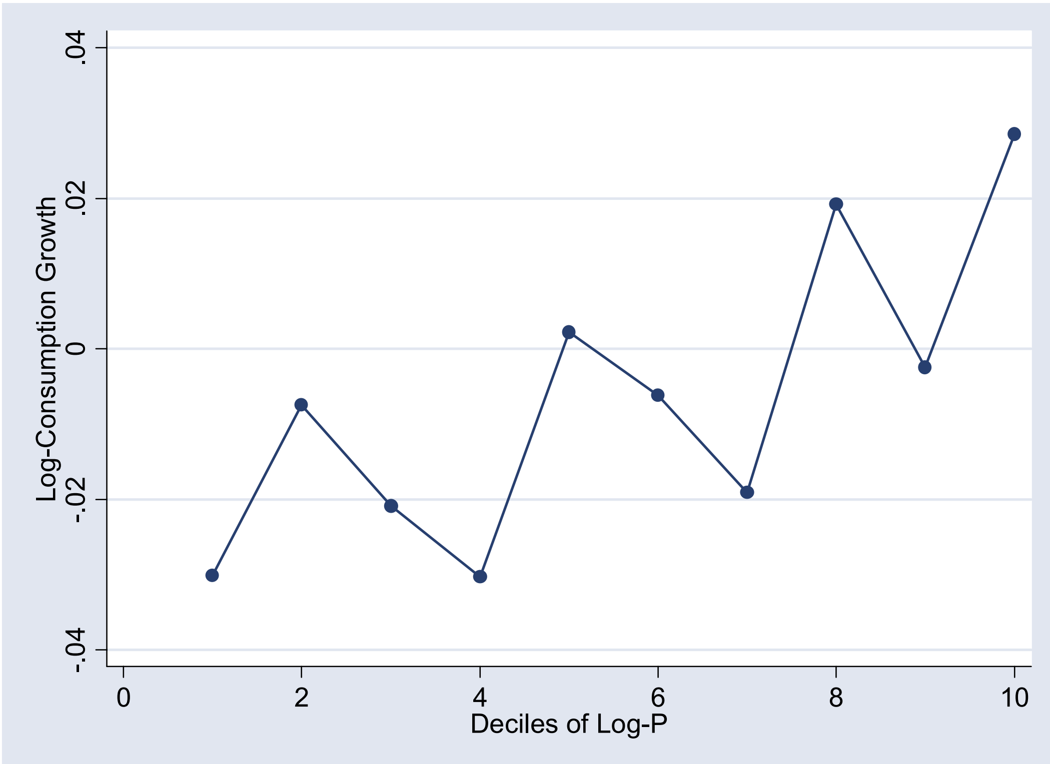

Figure 7: Mean Log-Consumption Growth by Decile of Log-Survival Probability (Singles Only)

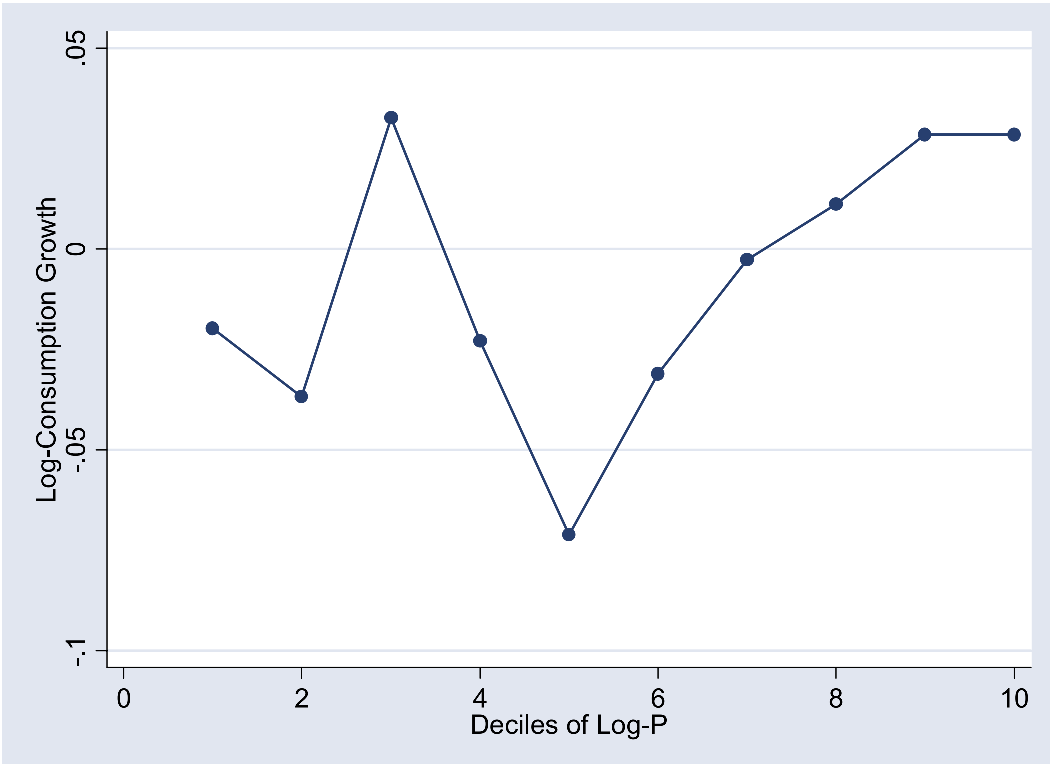

Figure 8: Mean Log-Consumption Growth by Decile of Log-Survival Probability (error bars= 1 standard deviation on either side of mean)

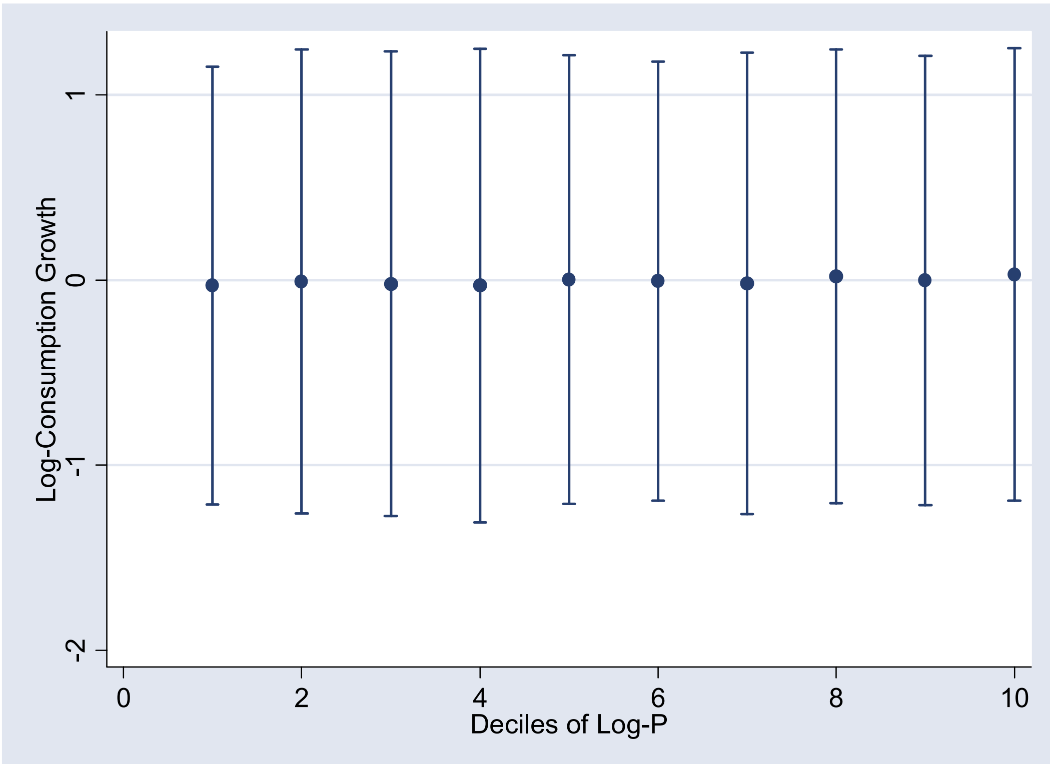

Figure 9: Mean Log-Consumption Growth by Decile of Log-Survival Probability (error bars= 1 standard deviation on either side of mean, singles only)